arXiv:1603.04547v1 [math.OC] 15 Mar 2016

Stochastic quasi-Newton methods for non-strongly convex problems: convergence and rate analysis

Farzad Yousefian[1] and Angelia Nedić[2] and Uday V. Shanbhag[3]

***Abstract*— Motivated by applications in optimization and machine learning, we consider stochastic quasi-Newton (SQN) methods for solving stochastic optimization problems. In the literature, the convergence analysis of these algorithms relies on strong convexity of the objective function. To our knowledge, no theoretical analysis is provided for the rate statements in the absence of this assumption. Motivated by this gap, we allow the objective function to be merely convex and we develop a cyclic regularized SQN method where the gradient mapping and the Hessian approximation matrix are both regularized at each iteration and are updated in a cyclic manner. We show that, under suitable assumptions on the stepsize and regularization parameters, the objective function value converges to the optimal objective function of the original problem in both almost sure and the expected senses. For each case, a class of feasible sequences that guarantees the convergence is provided. Moreover, the rate of convergence in terms of the objective function value is derived. Our empirical analysis on a binary classification problem shows that the proposed scheme performs well compared to both classic regularization SQN schemes and stochastic approximation method.**

I. Introduction

In this paper, we study a stochastic optimization problem of the form:

$$\min_{x\in\mathbb{R}^n} f(x) := \mathsf{E}[F(x,\xi(\omega))]\,, \tag{1}$$

where $F:\mathbb{R}^n\times\mathbb{R}^d\to\mathbb{R}$ is a function, the random vector ξ is defined as $\xi:\Omega\to\mathbb{R}^d$, $(\Omega,\mathcal{F},\mathbb{P})$ denotes the associated probability space and the expectation $\mathsf{E}[F(x,\xi)]$ is taken with respect to $\mathbb{P}$. A variety of applications can be cast as the model (1) (see [1], [2], [3], [4]). An important example in machine learning is the support vector machine (SVM) problem ([5], [6], [7]). In such problems, a training set containing a large number of input/output pairs $\{(u_i,v_i)\}_{i=1}^N\in\mathbb{R}^m\times\mathbb{R}$ is given where $v_i\in\{-1,1\}$ is the class' index. The goal is to learn a classifier (e.g., a hyperplane) $h(x,u)$ where x is the vector of parameters of the function h and u is the input data. To measure the distance of an observed output v_i from the classifier function h, a real-valued convex loss function $\ell(h;v)$ is defined. The objective function is considered as the following averaged loss over the training set

$$f(x) := \frac{1}{N}\sum_{i=1}^{N}\ell(h(x,u_i);v_i). \tag{2}$$

The preceding objective can be seen as a stochastic optimization model of the form (1), where $F(x,\xi) := \ell(h(x,u);v)$ and $\xi=(u,v)$.

Although, problem (1) may be seen as a deterministic problem, the challenges still arise when standard deterministic schemes are employed. In particular, when the expectation is over a general measure space (making computation of $\nabla_x\mathsf{E}[F(x,\xi)]$ difficult or impossible) or the distribution $\mathbb{P}$ is unavailable, standard gradient or Newton-based schemes cannot be directly applied. This has led to significant research on Monte-Carlo sampling techniques. Monte Carlo simulation methods have been used widely in the literature to solve stochastic optimization problems. Of these, sample average approximation (SAA) methods [8] and stochastic approximation (SA) methods ([9], [10]) (also referred to as stochastic gradient descent methods in the context of optimization) are among popular approaches. It has been discussed that when the sample size is large, the computational effort for implementing SAA schemes does not scale with the number of samples and these methods become inefficient ([10], [11]). SA methods, introduced by Robbins and Monro [9], require the construction of a sequence $\{x_k\}$, given a randomly generated $x_0\in\mathbb{R}^n$:

$$x_{k+1} := x_k - \gamma_k\nabla F(x_k,\xi_k), \quad \text{for } k\geq 0, \tag{SA}$$

where $\gamma_k>0$ denotes the stepsize and $\nabla F(x_k,\xi_k)$ denotes the sampled gradient of the function f with respect to x at x_k. Note that the gradient $\nabla F(x_k,\xi_k)$ is assumed to be an unbiased estimator of the true value of the gradient $\nabla f(x)$ at x_k, and assumed to be generated by a stochastic oracle. SA schemes are characterized by several disadvantages, including the poorer rate of convergence (than their deterministic counterparts) and the detrimental impact of conditioning on their performance. In deterministic regimes, when second derivatives are available, Newton schemes and their quasi-Newton counterparts have proved to be useful alternatives, particularly from the standpoint of displaying faster rates of convergence ([12], [13]).

Recently, there has been a growing interest in applying stochastic variants of quasi-Newton (SQN) methods for solving optimization and large scale machine learning problems. In these methods, x_k is given by the following update rule:

$$x_{k+1} := x_k - \gamma_k H_k^{-1}\nabla F(x_k,\xi_k), \quad \text{for } k\geq 0, \tag{SQN}$$

[1]Assistant Professor, School of Industrial Engineering & Management, Oklahoma State University, Stillwater, OK 74078, USA farzad.yousefian@okstate.edu

[2]Associate Professor, Industrial & Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA angelia@illinois.edu

[3]Associate Professor, Industrial & Manufacturing Engineering, Pennsylvania State University, University Park, PA 16802, USA udaybag@psu.edu

where $H_k \succeq 0$ is an approximation of the Hessian matrix at iteration k that incorporates the curvature information of the objective function within the algorithm. The convergence of this class of algorithms can be derived under a careful choice of the matrix H_k and the stepsize sequence γ_k. In particular, boundedness of the eigenvalues of H_k is an important factor in achieving global convergence in convex and nonconvex problems ([14], [15]). While in [16] the performance of SQN methods displayed to be favorable in solving high dimensional problems, Mokhtari et al. [17] developed a regularized BFGS method (RES) by updating the matrix H_k according to a modified version of BFGS update rule to assure convergence. To address large scale applications, limited memory variants (L-BFGS) were employed to ascertain scalability in terms of the number of variables ([6], [18]). In a recent extension [19], a stochastic quasi-Newton method is presented for solving nonconvex stochastic optimization problems. Also, a variance reduced SQN method with a constant stepsize was developed [20] for smooth strongly convex problems characterized by a linear convergence rate.

Motivation: One of the main assumptions in the developed stochastic SQN method (e.g. [6], [18]) is the strong convexity of the objective function. Specifically, this assumption plays an important role in deriving the rate of convergence of the algorithm. However, in many applications, the objective function is convex, but not strongly convex such as, for example, the logistic regression function that is given by $\ell(u^T x, v) := \ln(1+\exp(-u^T xv))$ for $u, x \in \mathbb{R}^n$ and $v \in \mathbb{R}$. While lack of strong convexity might lead to a very slow convergence, no theoretical results on the convergence rate in given in the literature of stochastic SQN methods. A simple remedy to address this challenge is to regularize the objective function with the term $\frac{1}{2}\mu\|x\|^2$ and solve the approximate problem of the form

$$\min_{x\in\mathbb{R}^n} f(x) + \frac{\mu}{2}\|x\|^2, \tag{3}$$

where $\mu > 0$ is the regularization parameter. A trivial drawback of this technique is that the optimal solution to the approximate problem (3) is not an optimal of the original problem (3). Importantly, choosing μ to be a small number deteriorates the convergence rate of the algorithm. This issue is resolved in SA schemes through employing averaging techniques for non-strongly convex problems and they display the optimal rate of $\mathcal{O}\left(\frac{1}{\sqrt{k}}\right)$ (see [10], [21]). A limitation to averaging SA schemes is that boundedness of the gradient mapping is required to achieve such a rate.

Contributions: Motivated by these gaps, in this paper, we consider stochastic optimization problems with non-strongly convex objective functions and Lipschitz but possibly unbounded gradient mappings. We develop a so-called cyclic regularized stochastic BFGS algorithm to solve this class of problems. Our framework is general and can be adapted within other variants of SQN methods. Unlike the classic regularization, we allow the regularization parameter μ, denoted by μ_k, to be updated and decay to zero through implementing the iterations. This enables the generated sequence to approach to the optimal solution of the original problem and also, benefits the scheme by guaranteeing a derived rate of convergence. A challenge in employing this technique is to maintain the secant condition and ascertain the positive definiteness of the BFGS matrix. We overcome this difficulty by carefully updating the regularization parameter and the BFGS matrix in a cyclic manner. We show that, under suitable assumptions on the stepsize and the regularization parameter (referred to as tuning sequences), the objective function value converges to the exact optimal value in an almost sure sense. Moreover, we show that under different settings, the algorithm achieves convergence in mean and we derive and upper bound for the error of the algorithm in terms of the tuning sequences. We complete our analysis by showing that under a specific choice of the tuning sequences, the rate of convergence in terms of the objective function value is of the order $\frac{1}{\sqrt[5]{k}}$.

The rest of the paper is organized as follows. Section II presents the outline of the proposed algorithm addressing problems with non-strongly convex objectives. In Section III, we prove the convergence of the scheme in both almost sure and expected senses and derive the rate statement. We present the numerical experiments in Section IV. The paper ends with some concluding remarks in Section V.

Notation: A vector x is assumed to be a column vector and x^T denotes its transpose, while $\|x\|$ denotes the Euclidean vector norm, i.e., $\|x\| = \sqrt{x^T x}$. We write *a.s.* as the abbreviation for "almost surely". For a symmetric matrix B, we write $\lambda_{min}(B)$ to denote its smallest eigenvalue. We use $\mathsf{E}[z]$ to denote the expectation of a random variable z. A function $f : X \subset \mathbb{R}^n \to \mathbb{R}$ is said to be strongly convex with parameter $\mu > 0$, if $f(y) \geq f(x) + \nabla f(y)^T(y-x) + \frac{\mu}{2}\|x-y\|^2$, for any $x, y \in X$. A mapping $F : X \subset \mathbb{R}^n \to \mathbb{R}$ is Lipschitz continuous with parameter $L > 0$ if for any $x, y \in X$, we have $\|F(x) - F(y)\| \leq L\|x-y\|$.

II. Outline of the algorithm

We begin by stating our general assumptions for problem (1). The underlying assumption in this paper is that the function f is convex and smooth.

Assumption 1: (a) The function $F(x,\xi)$ is convex with respect to x for any $\xi \in \Omega$.

(b) $f(x)$ is continuously differentiable with Lipschitz continuous gradients over $\mathbb{R}^n$ with parameter $L > 0$.

(c) The optimal solution set of problem (1) is nonempty.

Next, we state the assumptions on the random variable ξ and the properties of the stochastic estimator of the gradient mapping, i.e. ∇F.

Assumption 2: (a) Random variables ξ_k are i.i.d. for any $k \geq 0$;

(b) The stochastic gradient mapping $\nabla F(x,\xi)$ is an unbiased estimator of $\nabla f(x)$, i.e. $\mathsf{E}[\nabla F(x,\xi)] = \nabla f(x)$, and has bounded variance, i.e., there exists a scalar $\nu > 0$ such that $\mathsf{E}\big[\|\nabla F(x,\xi) - \nabla f(x)\|^2\big] \leq \nu^2$ for any $x \in \mathbb{R}^n$.

To solve (1), we propose a regularization algorithm that

generates a sequence $\{x_k\}$ for any $k \geq 0$:

$$x_{k+1} := x_k - \gamma_k \left(B_k^{-1} + \delta_k \mathbf{I}\right) \left(\nabla F(x_k, \xi_k) + \mu_k x_k\right). \tag{CR-SQN}$$

Here, $\gamma_k > 0$ denotes the stepsize at iteration k, B_k denotes the approximation of the Hessian matrix, $\mu_k > 0$ is the regularization parameter of the gradient mapping where

$$\begin{cases} \mu_{k+1} = \mu_k, & \text{if } k \text{ is even} \\ \mu_{k+1} < \mu_k, & \text{if } k \text{ is odd,} \end{cases} \tag{4}$$

while $\delta_k > 0$ is the regularization parameter of the matrix H_k. We assume when k is even, $\mu_k < \mu_{k-1}$ is chosen such that $\nabla F(x_k, \xi_k) + \mu_k x_k \neq 0$.[1] Let us define the matrix B_k by the following rule:

$$B_{k+1} := \begin{cases} B_k - \frac{B_k s_k s_k^T B_k}{s_k^T B_k s_k} + \frac{y_k^B (y_k^B)^T}{s_k^T y_k^B} + \rho \mu_k \mathbf{I}, & k \text{ even} \\ B_k, & k \text{ odd,} \end{cases} \tag{5}$$

where for an even k,

$$s_k := x_{k+1} - x_k,$$
$$y_k^B := \nabla F(x_{k+1}, \xi_k) - \nabla F(x_k, \xi_k) + (1-\rho)\mu_k s_k,$$

and $0 < \rho < 1$ is the regularization factor of the matrix B_{k+1} at iteration k. To state the properties of the matrix B_k, we start by defining the regularized function.

Definition 1: Consider the sequence $\{\mu_k\}$ of positive scalars. The regularized function $f_k : \mathbb{R}^n \to \mathbb{R}$ is defined as follows:

$$f_k(x) := f(x) + \frac{\mu_k}{2}\|x\|^2, \quad \text{for any } k \geq 0.$$

Similar notation can be used for the regularized stochastic function F_k as $F_k(x, \xi) := F(x, \xi) + \frac{\mu_k}{2}\|x\|^2$. We can now define the term y_k^{reg} as the difference between the value of the regularized stochastic gradient mappings at two consecutive points as follows:

$$y_k^{reg} := \nabla F_k(x_{k+1}, \xi_k) - \nabla F_k(x_k, \xi_k).$$

In the following result, we show that at iterations that the matrix B_k is updated, the secant condition is satisfied implying that B_k is well-defined. Also, we show that B_k is positive definite for $k \geq 0$.

Lemma 1: Let Assumption 1(a) hold, and let B_k be given by the update rule (5). Suppose $B_0 \succeq \rho\mu_0 \mathbf{I}$ is a symmetric matrix. Then, for any even k, the secant condition holds, i.e., $s_k^T y_k^B > 0$, and $B_{k+1} s_k = y_k^{reg}$. Moreover, for any k, B_k is symmetric and $B_{k+1} \succeq \rho\mu_k \mathbf{I}$.

Proof: It can be easily seen, by the induction on k, that all B_k are symmetric when B_0 is symmetric, assuming that the matrices are well defined. We use the induction on even values of k to show that the other statements hold and that the matrices are well defined. Suppose $k \geq 2$ is even and for any even values of t with $t < k$, we have $s_t^T y_t^B > 0$, $B_{t+1} s_t = y_t^{reg}$, and $B_{t+1} \succeq \rho\mu_t \mathbf{I}$. We show that all of these relations hold for $t = k$, i.e., $s_k^T y_k^B > 0$, $B_{k+1} s_k = y_k^{reg}$, and $B_{k+1} \succeq \rho\mu_k \mathbf{I}$. First, we prove that the secant condition holds. We can write

$$\begin{aligned} s_k^T y_k^B &= (x_{k+1} - x_k)^T (\nabla F(x_{k+1}, \xi_k) - \nabla F(x_k, \xi_k) \\ &\quad + (1-\rho)\mu_k (x_{k+1} - x_k)) \\ &\geq (1-\rho)\mu_k \|x_{k+1} - x_k\|, \end{aligned}$$

where we used the convexity of $F(\cdot, \xi)$.

From the induction hypothesis, $B_{k-1} \succeq \rho\mu_{k-2}\mathbf{I}$ since $k-2$ is even. Furthermore, since $k-1$ is odd, we have $B_k = B_{k-1}$ by the update rule (5). Therefore, B_k is positive definite. Note that since k is even, the choice of μ_k is such that $\nabla F(x_k, \xi_k) + \mu_k x_k \neq 0$ (see the discussion following (4)). Since B_k is positive definite, the matrix $B_k^{-1} + \delta_k \mathbf{I}$ is positive definite. Therefore, we have $\left(B_k^{-1} + \delta_k \mathbf{I}\right)\left(\nabla F(x_k, \xi_k) + \mu_k x_k\right) \neq 0$, implying that $x_{k+1} \neq x_k$. Hence

$$s_k^T y_k^B \geq (1-\rho)\mu_k \|x_{k+1} - x_k\|^2 > 0,$$

where we used $\rho < 1$. Thus, the secant condition holds.

Also, since $k-1$ is odd, by update rule (4), it follows that $\mu_k = \mu_{k-1}$. From the update rule (5) and that B_k is positive definite and symmetric, we have

$$B_{k+1} \succeq B^{0.5}\left(\mathbf{I} - \frac{B_k^{0.5} s_k (s_k)^T B_k^{0.5}}{\|B_k^{0.5} s_k\|^2}\right) B^{0.5} + \rho\mu_k \mathbf{I}, \tag{6}$$

where the last relation is due to $y_k^B (y_k^B)^T \succeq 0$ and $s_k^T y_k^B > 0$.

Since $x_{k+1} \neq x_k$ we have $s_k^T B_k s_k \neq 0$. Thus, the matrix $\mathbf{I} - \frac{B_k^{0.5} s_k (s_k)^T B_k^{0.5}}{\|s_k^T B_k^{0.5}\|^2}$ is well defined and positive semidefinite, since it is symmetric with eigenvalues are between 0 and 1.

Next, we show that B_{k+1} satisfies $B_{k+1} s_k = y_k^{reg}$. Using the update rule (5), for even k we have,

$$\begin{aligned} B_{k+1} s_k &= B_k s_k - \frac{B_k s_k s_k^T B_k s_k}{s_k^T B_k s_k} + \frac{y_k^B (y_k^B)^T s_k}{s_k^T y_k^B} + \rho\mu_k s_k \\ &= B_k s_k - B_k s_k + y_k^B + \rho\mu_k s_k \\ &= \nabla F(x_{k+1}, \xi_k) - \nabla F(x_k, \xi_k) + \mu_k s_k. \end{aligned}$$

Since k is even, we have $\mu_{k+1} = \mu_k$ implying that

$$\begin{aligned} B_{k+1} s_k &= \nabla F(x_{k+1}, \xi_k) - \nabla F(x_k, \xi_k) \\ &\quad + \mu_{k+1} x_{k+1} - \mu_k x_k \\ &= \nabla F_k(x_{k+1}, \xi_k) - \nabla F_k(x_k, \xi_k) = y_k^{reg}, \end{aligned}$$

where the last equality follows by the definition of the regularized mappings.

From the preceding discussion, we conclude that the induction hypothesis holds also for $t = k$. Therefore, all the desired results hold for any even k. To complete the proof, we need to show that for any odd k, we have $B_{k+1} = B_k \succeq \rho\mu_k \mathbf{I}$. By the update rule (5), we have $B_{k+1} = B_k$. Since $k-1$ is even, $B_k \succeq \rho\mu_{k-1}\mathbf{I}$. Also, from (4) we have $\mu_k = \mu_{k-1}$. Therefore, $B_{k+1} = B_k \succeq \rho\mu_k \mathbf{I}$. ■

[1] Note that we can ensure this non-zero condition, as follows. If we have $\nabla F(x_k, \xi_k) + \mu_k x_k = 0$ and $x_k \neq 0$, then by replacing μ_k with a smaller value the relation will hold. If $x_k = 0$, then we can draw a new sample of ξ_k to satisfy the relation. We can get stuck at x_k with sampling if 0 solves the original problem.

III. Convergence Analysis

In this section, we analyze the convergence properties of the stochastic recursion (CR-SQN). The following assumption provides the required conditions on the stepsize sequence γ_k and is a commonly used assumption in the regime of stochastic approximation methods [17], [19], [6].

Property 1: [Properties of the regularized function] The function f_k from Definition 1 for any $k \geq 0$ has the following properties:

(a) f_k is strongly convex with a parameter μ_k.
(b) f_k has Lipschitzian gradients with parameter $L + \mu_k$.
(c) f_k has a unique minimizer over $\mathbb{R}^n$, denoted by x_k^*. Moreover, for any $x \in \mathbb{R}^n$,

$$2\mu_k(f_k(x) - f_k(x_k^*)) \leq \|\nabla f_k(x)\|^2,$$

$$\|\nabla f_k(x)\|^2 \leq 2(L + \mu_k)(f_k(x) - f_k(x_k^*)).$$

The existence and uniqueness of x_k^* in Property 1(c) is due to the strong convexity of the function f_k (see, for example, Sec. 1.3.2 in [22]), while the relation for the gradient is known to hold for a strongly convex function with a parameter μ that also has Lipschitz gradients with a parameter L (see Lemma 1 in page 23 in [22]).

The next result provides an important property for the recursion (CR-SQN) that will be subsequently used to show the convergence of the scheme. Throughout, we let $\mathcal{F}_k$ denote the history of the method up to time k, i.e., $\mathcal{F}_k = \{x_0, \xi_0, \xi_1, \ldots, \xi_{k-1}\}$ for $k \geq 1$ and $\mathcal{F}_0 = \{x_0\}$. Also, we denote the stochastic error of the regularized gradient estimator by

$$w_k := \nabla F(x_k, \xi_k) - \nabla f(x_k), \quad \text{for all } k \geq 0. \tag{7}$$

Lemma 2: [A recursive error bound inequality] Consider the algorithm (CR-SQN). Suppose sequences γ_k, δ_k, and μ_k are chosen such that for any $k \geq 0$, μ_k satisfies (4), and

$$(L + \mu_k)^2 \gamma_k \left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2 \leq \delta_k \mu_k. \tag{8}$$

Under Assumptions 1 and 2, for any $k \geq 1$ and any optimal solution x^*, we have

$$\begin{aligned} &\mathsf{E}[f_{k+1}(x_{k+1}) \mid \mathcal{F}_k] - f^* \\ &\leq (1 - \gamma_k\delta_k\mu_k)\left(f_k(x_k) - f^*\right) + \gamma_k\delta_k\frac{\mu_k^2}{2}\|x^*\|^2 \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2. \end{aligned} \tag{9}$$

Proof: The Lipschitzian property of ∇f_k (see Property 1(b)) and the recursion (CR-SQN) imply that

$$\begin{aligned} &f_k(x_{k+1}) \leq f_k(x_k) + \nabla f_k(x_k)^T(x_{k+1} - x_k) \\ &+ \frac{(L + \mu_k)}{2}\|x_{k+1} - x_k\|^2 \\ &= f_k(x_k) - \gamma_k\nabla f_k(x_k)^T\left(B_k^{-1} + \delta_k\mathbf{I}\right)\left(\nabla F(x_k, \xi_k) + \mu_k x_k\right) \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\|\left(B_k^{-1} + \delta_k\mathbf{I}\right)\left(\nabla F(x_k, \xi_k) + \mu_k x_k\right)\|^2. \end{aligned}$$

From the definition of the stochastic error w_k (see (7)) and and the definition of the regularized function (see Definition 1), we have

$$\nabla F(x_k, \xi_k) + \mu_k x_k = \nabla f(x_k) + \mu_k x_k + w_k = \nabla f_k(x_k) + w_k.$$

Hence,

$$\begin{aligned} &f_k(x_{k+1}) \leq f_k(x_k) \\ &- \gamma_k\nabla f_k(x_k)^T(B_k^{-1} + \delta_k\mathbf{I})(\nabla f_k(x_k) + w_k) \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\|\left(B_k^{-1} + \delta_k\mathbf{I}\right)\left(\nabla f_k(x_k) + w_k\right)\|^2 \\ &\leq f_k(x_k) - \gamma_k\lambda_{min}\left(B_k^{-1} + \delta_k\mathbf{I}\right)\|\nabla f_k(x_k)\|^2 \\ &-\gamma_k\nabla f_k(x_k)^T\left(B_k^{-1} + \delta_k\mathbf{I}\right)w_k \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\|B_k^{-1} + \delta_k\mathbf{I}\|^2\left(\|\nabla f_k(x_k) + w_k\|^2\right). \end{aligned}$$

Note that Lemma 1(b) implies that

$$\delta_k\mathbf{I} \preceq B_k^{-1} + \delta_k\mathbf{I} \preceq \left((\rho\mu_{k-1})^{-1} + \delta_k\right)\mathbf{I}.$$

From the preceding two relations we obtain

$$\begin{aligned} &f_k(x_{k+1}) \leq f_k(x_k) - \gamma_k\delta_k\|\nabla f_k(x_k)\|^2 \\ &- \gamma_k\nabla f_k(x_k)^T\left(B_k^{-1} + \delta_k\mathbf{I}\right)w_k \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\|\nabla f_k(x_k) + w_k\|^2. \end{aligned}$$

Next, we take the expected value from both sides with respect to $\mathcal{F}_k$. Note that the matrix B_k and x_k are both deterministic parameters if the history $\mathcal{F}_k$ is known. Note that from Assumption 2, $\mathsf{E}[w_k \mid \mathcal{F}_k] = 0$ and $\mathsf{E}\left[\|w_k\|^2 \mid \mathcal{F}\right] \leq \nu^2$. Therefore,

$$\mathsf{E}\left[\|\nabla f_k(x_k) + w_k\|^2 \mid \mathcal{F}_k\right] \leq \|\nabla f_k(x_k)\|^2 + \nu^2.$$

Thus, we obtain

$$\begin{aligned} &\mathsf{E}[f_k(x_{k+1}) \mid \mathcal{F}_k] \leq f_k(x_k) - \gamma_k\delta_k\|\nabla f_k(x_k)\|^2 \\ &+ \frac{(L + \mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\left(\|\nabla f_k(x_k)\|^2 + \nu^2\right). \end{aligned}$$

Using Property 1(c) for the function f_k, we have

$$\begin{aligned} &\mathsf{E}[f_k(x_{k+1}) \mid \mathcal{F}_k] \leq f_k(x_k) - 2\gamma_k\delta_k\mu_k\left(f_k(x_k) - f_k(x_k^*)\right) \\ &+ (L + \mu_k)^2\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\left(f_k(x_k) - f_k(x_k^*)\right) \\ &+ \frac{(L + \mu_k)^2}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2. \end{aligned}$$

Due to our assumption on the choice of the sequences γ_k, δ_k and μ_k, we have $(L + \mu_k)^2\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2 \leq \gamma_k\delta_k\mu_k$. Thus,

$$\begin{aligned} &\mathsf{E}[f_k(x_{k+1}) \mid \mathcal{F}_k] \leq f_k(x_k) - \gamma_k\delta_k\mu_k\left(f_k(x_k) - f_k(x_k^*)\right) \\ &+ \frac{(L + \mu_k)^2}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2. \end{aligned} \tag{10}$$

In the last step, we build a recursive inequality for the error term $f_k(x_k) - f^*$. Adding and subtracting f^*, we obtain

$$\begin{aligned} f_k(x_k) - f_k(x_k^*) &= \left(f_k(x_k) - f^*\right) + \left(f^* - f_k(x_k^*)\right) \\ &= \left(f_k(x_k) - f^*\right) + \left(f_k(x^*) - \frac{\mu_k}{2}\|x^*\|^2 - f_k(x_k^*)\right), \end{aligned}$$

where the last equality follows from $f^* = f_k(x^*) - \frac{\mu_k}{2}\|x^*\|^2$. Since x_k^* is the minimizer of $f_k(x_k^*)$, we have $f_k(x^*) - f_k(x_k^*) \geq 0$, implying that

$$f_k(x_k) - f_k(x_k^*) \geq f_k(x_k) - f^* - \frac{\mu_k}{2}\|x^*\|^2.$$

By substituting the preceding inequality in (10), we obtain

$$\mathsf{E}[f_k(x_{k+1}) \mid \mathcal{F}_k] \leq f_k(x_k) - \gamma_k\delta_k\mu_k\left(f_k(x_k) - f^*\right) + \gamma_k\delta_k\frac{\mu_k^2}{2}\|x^*\|^2 + \frac{(L+\mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2.$$

By subtracting f^* from both sides of the preceding inequality, we see that

$$\mathsf{E}[f_k(x_{k+1}) \mid \mathcal{F}_k] - f^* \leq (1 - \gamma_k\delta_k\mu_k)\left(f_k(x_k) - f^*\right) + \gamma_k\delta_k\frac{\mu_k^2}{2}\|x^*\|^2 + \frac{(L+\mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2.$$

Next, we relate the values $f_{k+1}(x_{k+1})$ and $f_k(x_{k+1})$. From Definition 1 and μ_k being non-increasing we can write

$$\begin{aligned} f_{k+1}(x_{k+1}) &= f(x_{k+1}) + \frac{\mu_{k+1}}{2}\|x_{k+1}\|^2 \\ &\leq f(x_{k+1}) + \frac{\mu_k}{2}\|x_{k+1}\|^2 = f_k(x_{k+1}). \end{aligned}$$

Therefore, the desired inequality (9) holds. ∎

We make use of the following result, which can be found in [22] (see Lemma 11 on page 50).

Lemma 3: Let $\{v_k\}$ be a sequence of nonnegative random variables, where $\mathsf{E}[v_0] < \infty$, and let $\{\alpha_k\}$ and $\{\beta_k\}$ be deterministic scalar sequences such that:

$$\mathsf{E}[v_{k+1}|v_0,\ldots,v_k] \leq (1-\alpha_k)v_k + \beta_k \quad a.s. \text{ for all } k \geq 0,$$
$$0 \leq \alpha_k \leq 1, \quad \beta_k \geq 0,$$
$$\sum_{k=0}^{\infty}\alpha_k = \infty, \quad \sum_{k=0}^{\infty}\beta_k < \infty, \quad \lim_{k\to\infty}\frac{\beta_k}{\alpha_k} = 0.$$

Then, $v_k \to 0$ almost surely.

In order to apply Lemma 3 to the inequality (9) and prove the almost sure convergence, we use the following definitions:

$$v_k := f_k(x_k) - f^*, \quad \alpha_k := \gamma_k\delta_k\mu_k,$$
$$\beta_k := \gamma_k\delta_k\mu_k^2\|x^*\|^2 + (L+\mu_k)\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2. \tag{11}$$

To satisfy the conditions of Lemma 3, we identify a set of sufficient conditions on the sequences $\{\gamma_k\}, \{\mu_k\}$, and $\{\delta_k\}$ in forthcoming assumption. Later in Lemma 4, we provide a class of sequences that meet these assumptions.

Assumption 3: [Sufficient conditions on sequences for a.s. convergence] Let the sequences $\{\gamma_k\}, \{\mu_k\}$, and $\{\delta_k\}$ be non-negative and satisfy the following conditions:

(a) $\lim_{k\to\infty}\frac{\gamma_k}{\mu_k^3\delta_k} = 0$;
(b) $\delta_k\mu_{k-1} \leq 1$ for $k \geq 1$;
(c) μ_k satisfies (4) and $\mu_k \to 0$;
(d) $\sum_{k=0}^{\infty}\gamma_k\delta_k\mu_k = \infty$;
(e) $\sum_{k=0}^{\infty}\left(\frac{\gamma_k}{\mu_k}\right)^2 < \infty$;
(f) $\sum_{k=0}^{\infty}\gamma_k\delta_k\mu_k^2 < \infty$;
(g) $\gamma_k\delta_k\mu_k \leq 1$ for $k \geq 0$;

With Assumption 3, we have the following result.

Theorem 1: [Almost sure convergence] Consider the algorithm (CR-SQN). Suppose Assumptions 1, 2 and 3 hold. Then, $\lim_{k\to\infty} f(x_k) = f^*$ a.s.

Proof: First, note that from Assumption 3(a), there exists $K \geq 1$ such that for any $k \geq K$ we have $\frac{4\gamma_k}{\rho^2\mu_k^3\delta_k} \leq \frac{1}{(L+\mu_0)^2}$. Taking this into account, using Assumption 3(b) and (c) we can write

$$\begin{aligned} \gamma_k\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2 &\leq \gamma_k\left(2(\rho\mu_{k-1})^{-1}\right)^2 \\ &\leq \frac{4\gamma_k}{\rho^2\mu_{k-1}^2} \leq \frac{4\gamma_k}{\rho^2\mu_k^2} = (\mu_k\delta_k)\frac{4\gamma_k}{\rho^2\delta_k\mu_k^3} \leq \frac{\delta_k\mu_k}{(L+\mu_k)^2}, \end{aligned}$$

implying that condition (8) of Lemma 2 holds. Hence, relation (9) holds for any $k \geq K$. Next, we apply Lemma 3 to prove a.s. convergence of the algorithm (CR-SQN). Consider the definitions in (11) for any $k \geq K$. The non-negativity of α_k and β_k is implied by the definition and that γ_k, δ_k, and μ_k are positive. From (9), we have

$$\mathsf{E}[v_{k+1} \mid \mathcal{F}_k] \leq (1-\alpha_k)v_k + \beta_k \quad \text{for all } k \geq K.$$

Since $f^* \leq f(x)$ for any arbitrary $x \in \mathbb{R}^n$, we can write

$$v_k = f_k(x_k) - f^* = (f(x_k) - f^*) + \frac{\mu_k}{2}\|x_k\|^2 \geq 0.$$

From Assumption 3(g), we obtain $\alpha_k \leq 1$. Also, from Assumption 3(d), we get $\sum_{k=K}^{\infty}\alpha_k = \infty$. Using Assumption 3(b) and the definition of β_k in (11), for an arbitrary solution x^*, we can write

$$\sum_{k=K}^{\infty}\beta_k \leq \|x^*\|^2\sum_{k=K}^{\infty}\gamma_k\delta_k\mu_k^2 + 4(L+\mu_0)\frac{\nu^2}{\rho^2}\sum_{k=0}^{\infty}\frac{\gamma_k^2}{\mu_k^2} < \infty,$$

where the last inequality is deduced by Assumptions 3(e) and 3(f). Similarly, we can write

$$\begin{aligned} \lim_{k\to\infty}\frac{\beta_k}{\alpha_k} &\leq \|x^*\|^2\lim_{k\to\infty}\mu_k + 4(L+\mu_0)\frac{\nu^2}{\rho^2}\lim_{k\to\infty}\frac{\gamma_k^2\mu_k^{-2}}{\gamma_k\delta_k\mu_k} \\ &= \|x^*\|^2\lim_{k\to\infty}\mu_k + 4(L+\mu_0)\frac{\nu^2}{\rho^2}\lim_{k\to\infty}\frac{\gamma_k}{\mu_k^3\delta_k} = 0, \end{aligned}$$

where the last equation is implied by Assumptions3(a) and 3(c). Therefore, all the conditions of Lemma 3 hold and we conclude that $v_k := f_k(x_k) - f^*$ converges to 0 a.s. Let us define $v_k' := f(x_k) - f^*$ and $v_k'' := \frac{\mu_k}{2}\|x_k\|^2$, so that $v_k = v_k' + v_k''$. Since v_k' and v_k'' are non-negative, and $v_k \to 0$ a.s., it follows that $v_k' \to 0$ and $v_k'' \to 0$ a.s., implying that $\lim_{k\to\infty} f(x_k) = f^*$ a.s. ∎

Lemma 4: Let the sequences γ_k, δ_k, and μ_k be given by the following rules:

$$\gamma_k = \frac{\gamma_0}{(k+1)^a}, \quad \delta_k = \frac{\delta_0}{(k+1)^b}, \quad \mu_k = \frac{\mu_0 2^c}{(k+\kappa)^c}, \tag{12}$$

where $\kappa = 2$ if k is even and $\kappa = 1$ otherwise, γ_k, δ_0, μ_0 are positive scalars such that $\delta_0\mu_0 \leq 2^b$ and $\gamma_0\delta_0\mu_0 \leq 1$, and a, b, and c are positive scalars that satisfy the following conditions:

$$a > 3c + b, \; a + b + c \leq 1$$
$$a - c > 0.5, \; a + 2c + b > 1.$$

Then, the sequences γ_k, δ_k, and μ_k satisfy Assumption 3.

Proof: In the following, we show that the presented class of sequences satisfy each of the conditions listed in Assumption 3:

(a) Replacing the sequences by their given rules we obtain

$$\frac{\gamma_k}{\mu_k^3\delta_k} = \frac{\gamma_0}{8^c\mu_0^3\delta_0}(k+1)^{-a+b}(k+\kappa)^{3c} \leq \frac{\gamma_0}{8^c\mu_0^3\delta_0}(k+1)^{-a+b+3c}.$$

Since $a > b + 3c$, the preceding term goes to zero verfying Assumption 3(a).

(b) The given rules (12) imply that δ_k and μ_k are both non-increasing sequences. Therefore, we have $\delta_k\mu_{k-1} \leq \delta_1\mu_0$ for any $k \geq 1$. So, to show that Assumption 3(b) holds, it is enought to show that $\delta_1\mu_0 \leq 1$. From (12) we have $\delta_1 = \delta_0 2^{-b}$. Since we assumed that $\delta_0\mu_0 \leq 2^b$, we can conclude that $\delta_1\mu_0 \leq 1$ implying that Assumption 3(b) holds.

(c) Let k be an even number. Thus, $\kappa = 2$. From (12) we have $\mu_k = \mu_{k+1} = \frac{\mu_0 2^c}{(k+2)^c}$. Now, let k be an odd number. Again, according to (12) can write

$$\mu_{k+1} = \frac{\mu_0 2^c}{((k+1)+2)^c} < \frac{\mu_0 2^c}{(k+1)^c} = \frac{\mu_0 2^c}{(k+\kappa)^c} = \mu_k.$$

Therefore, μ_k given by (12) satisfies (4). Also, from (12) we have $\mu_k \to 0$. Thus, Assumption 3(c) holds. μ_k satisfies (4) and $\mu_k \to 0$.

(d) From (12), we can write

$$\sum_{k=0}^{\infty}\gamma_k\delta_k\mu_k = \gamma_0\delta_0\mu_0 2^c\sum_{k=0}^{\infty}(k+1)^{-(a+b+c)} = \infty,$$

where the last inequality is due to the assumption that $a + b + c \leq 1$. Therefore, Assumption 3(d) holds.

(e) From (12), we have

$$\sum_{k=0}^{\infty}\left(\frac{\gamma_k}{\mu_k}\right)^2 = \frac{\gamma_0^2}{\mu_0^2 4^c}\sum_{k=0}^{\infty}\frac{(k+\kappa)^{2c}}{(k+1)^{2a}} \leq \frac{\gamma_0^2}{\mu_0^2 4^c}\left(\sum_{k=0}^{1}\frac{(k+\kappa)^{2c}}{(k+1)^{2a}} + \sum_{k=2}^{\infty}\frac{(2k)^{2c}}{k^{2a}}\right) < \infty$$

where the last inequality is due to $a - c > 0.5$. Therefore, Assumption 3(e) is verified.

(f) Using (12), it follows

$$\sum_{k=0}^{\infty}\gamma_k\delta_k\mu_k^2 = \gamma_0\delta_0\mu_0^2 4^c\sum_{k=0}^{\infty}(k+\kappa)^{-2c}(k+1)^{a+b} \leq \gamma_0\delta_0\mu_0^2 4^c\sum_{k=0}^{\infty}(k+1)^{-(2c+a+b)} < \infty,$$

where the last inequality is due to $a+2c+b > 1$. Therefore, Assumption 3(e) holds.

(g) The rules in (12) imply that γ_k ,δ_k and μ_k are all non-increasing sequences. We also assumed that $\gamma_0\delta_0\mu_0 \leq 1$. Hence, $\gamma_k\delta_k\mu_k \leq 1$ for any $k \geq 1$ and Assumption 3(f) holds. ∎

Remark 1: When $a = 0.75$, $b = 0$, and $c = 0.24$, and $\gamma_0 = \delta_0 = \mu_0 = 0.9$, the Assumption 3 is satisfied.

Assumption 4: [Sufficient conditions on sequences for convergence in mean] Let the sequences $\{\gamma_k\}, \{\mu_k\}$, and $\{\delta_k\}$ be non-negative and satisfy the following conditions:

(a) $\lim_{k\to\infty}\frac{\gamma_k}{\mu_k^3\delta_k} = 0$;
(b) $\delta_k\mu_{k-1} \leq 1$ for $k \geq 1$;
(c) μ_k satisfies (4);
(d) There exist $0 < \alpha < 1$ and $K_1 \geq 0$ such that $\frac{\gamma_{k-1}}{\mu_{k-1}^3\delta_{k-1}} \leq \frac{\gamma_k}{\mu_k^3\delta_k}(1 + \alpha\gamma_k\delta_k\mu_k)$ for $k \geq K_1$;
(e) There exist a scalar $\mathcal{B} > 0$ and $K_2 \geq 0$ such that $\delta_k\mu_k^4 \leq \mathcal{B}\gamma_k$ for $k \geq K_2$;

Theorem 2: [Convergence in mean] Consider the algorithm (CR-SQN). Suppose Assumptions 1, 2 and 4 hold. Then, there exists some $K \geq 0$ such that

$$\mathsf{E}[f(x_{k+1})] - f^* \leq \theta\frac{\gamma_k}{\mu_k^3\delta_k} \quad \text{for any } k \geq K, \tag{13}$$

where f^* is the optimal value of problem (1),

$$\theta = \max\left\{\frac{\mu_K^3\delta_K e_{K+1}}{\gamma_K}, \frac{0.5\mathcal{B}\|x^*\|^2 + 2(L+\mu_0)\frac{\nu^2}{\rho^2}}{1-\alpha}\right\},$$

with $e_{K+1} := \mathsf{E}[f_{K+1}(x_{K+1})] - f^*$.

Proof: Similar to the proof of Theorem 1, from Assumption 4(a), there exists $K \geq 1$ such that for any $k \geq K_0$, the condition (8) holds and, therefore, the inequality (9) holds. Let $K = \max\{K_0, K_1, K_2\}$. Taking expectation from both sides of (9), we obtain for any solution x^*,

$$e_{k+1} \leq (1 - \gamma_k\delta_k\mu_k)e_k + \gamma_k\delta_k\frac{\mu_k^2}{2}\|x^*\|^2 + \frac{(L+\mu_k)}{2}\gamma_k^2\left((\rho\mu_{k-1})^{-1} + \delta_k\right)^2\nu^2,$$

where $e_k := \mathsf{E}[f_k(x_k)] - f^*$. Using Assumption 4(b), (c) and (e), the preceding inequality yields

$$e_{k+1} \leq (1 - \gamma_k\delta_k\mu_k)e_k + \frac{\mathcal{B}}{2}\|x^*\|^2\frac{\gamma_k^2}{\mu_k^2} + 2(L+\mu_0)\frac{\nu^2}{\rho^2}\frac{\gamma_k^2}{\mu_k^2} = (1 - \gamma_k\delta_k\mu_k)e_k + \mathcal{B}_1\frac{\gamma_k^2}{\mu_k^2}, \qquad \text{for } k \geq K, \tag{14}$$

where $\mathcal{B}_1 := \frac{\mathcal{B}}{2}\|x^*\|^2 + 2(L+\mu_0)\frac{\nu^2}{\rho^2}$. We use induction to show the desired result. First, we show that (13) holds for $k = K$. We have

$$e_{K+1} = \mathsf{E}[f_{K+1}(x_{K+1})] - f^* = \left(\frac{\mu_K^3\delta_K(\mathsf{E}[f_{K+1}(x_{K+1})] - f^*)}{\gamma_K}\right)\frac{\gamma_K}{\mu_K^3\delta_K} \leq \theta\frac{\gamma_K}{\mu_K^3\delta_K},$$

implying that (13) holds for $k = K$. Now assume that $e_k \leq \theta\frac{\gamma_{k-1}}{\mu_{k-1}^3\delta_{k-1}}$, for some $k \geq K$. We show that $e_{k+1} \leq \theta\frac{\gamma_k}{\mu_k^3\delta_k}$. From the induction hypothesis and (14) we have

$$e_{k+1} \leq (1 - \gamma_k\delta_k\mu_k)\ \theta\frac{\gamma_{k-1}}{\mu_{k-1}^3\delta_{k-1}} + \mathcal{B}_1\frac{\gamma_k^2}{\mu_k^2},$$

Using Assumption 4(d) we obtain

$$e_{k+1} \leq (1 - \gamma_k\delta_k\mu_k)\theta\frac{\gamma_k}{\mu_k^3\ \delta_k}(1 + \alpha\gamma_k\delta_k\mu_k) + \mathcal{B}_1\frac{\gamma_k^2}{\mu_k^2}.$$

The definition of θ and $\mathcal{B}_1$ imply that the term $\theta(1-\alpha)-\mathcal{B}_1$ is non-negative. It follows

$$e_{k+1} \leq \theta\frac{\gamma_k}{\mu_k^3\delta_k} - \theta(1-\alpha)\frac{\gamma_k^2}{\mu_k^2} + \mathcal{B}_1\frac{\gamma_k^2}{\mu_k^2}$$
$$= \theta\frac{\gamma_k}{\mu_k^3\delta_k} - (\theta(1-\alpha)-\mathcal{B}_1)\frac{\gamma_k^2}{\mu_k^2} \leq \theta\frac{\gamma_k}{\mu_k^3\delta_k},$$

This shows that the induction argument holds true. Also, we have $f_{k+1}(x_{k+1}) = f(x_{k+1}) + \frac{\mu_{k+1}}{2}\|x_{k+1}\|^2 \geq f(x_{k+1})$. Therefore, we conclude that (13) holds. ■

Lemma 5: Let the sequences γ_k, δ_k, and μ_k be given by (12), where γ_k, δ_0, μ_0 are positive scalars such that $\delta_0\mu_0 \leq 2^b$, and a, b, and c are positive scalars that satisfy the following conditions:

$$a > 3c+b,\ a+b<1,\ -a+4c+b \geq 0.$$

Then, the sequences γ_k, δ_k, and μ_k satisfy Assumption 4.

Proof: In the following, we verify the conditions of Assumption 4.

Conditions (a), (b), and (c): This is already shown in parts (a), (b) and (c) of the proof of Lemma 4 due to $a > 3c+b$, $c > 0$, and $\delta_0\mu_0 \leq 2^b$.

(d) It suffices to show there exist K_1 and $\alpha \in (0,1)$ such that for any $k \geq K_1$

$$\frac{\gamma_{k-1}}{\gamma_k}\frac{\mu_k^3}{\mu_{k-1}^3}\frac{\delta_k}{\delta_{k-1}} - 1 \leq \alpha\gamma_k\delta_k\mu_k. \tag{15}$$

From (12), we obtain

$$\frac{\gamma_{k-1}}{\gamma_k}\frac{\mu_k^3}{\mu_{k-1}^3}\frac{\delta_k}{\delta_{k-1}} - 1 \leq \frac{\gamma_{k-1}}{\gamma_k} - 1 = \left(1+\frac{1}{k}\right)^a - 1$$
$$= 1 + \frac{a}{k} + o(\frac{1}{k}) - 1 = \mathcal{O}\left(\frac{1}{k}\right),$$

where the first inequality is implied due to both μ_k and δ_k are non-increasing sequences, and in the second equation we used the Taylor's expansion of $\left(1+\frac{1}{k}\right)^a$. Therefore, since the right hand-side of the relation (15) is of the order $\frac{1}{k^{a+b+c}}$ and that $a+b+c<1$, the preceding inequality shows that such α and K_1 exist such that Assumption 4(d) holds.

(e) From (12), we have

$$\frac{\delta_k\mu_k^4}{\gamma_k} = \delta_0\mu_0 2^c(k+\kappa)^{-4c}(k+1)^{a-b} \leq \frac{\delta_0\mu_0 2^c}{(k+1)^{-a+b+4c}}.$$

Since we assumed $-a+4c+b \geq 0$, there exists $\mathcal{B} > 0$ such that Assumption 4(e) is satisfied. ■

Theorem 3: [Rate of convergence] Consider the algorithm (CR-SQN). Suppose Assumptions 1 and 2 are satisfied. Let the sequences γ_k, δ_k, and μ_k be given by (12) with $a = 0.8$, $b = 0$, and $c = 0.2$, and $\delta_0 = \mu_0 = 0.9$ and $\gamma_0 > 0$. Then,

(a) $\lim_{k\to\infty} f(x_k) = f^*$ almost surely, where f^* is the optimal value of problem (1).

(b) We have

$$\mathsf{E}[f(x_k)] - f^* = \mathcal{O}(\frac{1}{\sqrt[5]{k}}).$$

Proof: (a) The given values of a, b, c and δ_0 and μ_0 satisfy the conditions of Lemma 4. Therefore, all conditions of Theorem 1 are met, and the desired statement follows.

(b) The given values of a, b, c and δ_0 and μ_0 satisfy the conditions of Lemma 5. Therefore, all conditions of Theorem 2 are satisfied, so from (13) we obtain

$$\mathsf{E}[f(x_{k+1})] - f^* \leq \theta\frac{\gamma_k}{\mu_k^3\delta_k} = \mathcal{O}\left(\frac{k^{-0.8}}{k^{-0.6}}\right) = \mathcal{O}(\frac{1}{\sqrt[5]{k}}).$$

■

Remark 2 (Computational cost): In large scale settings, a natural concern related to the implementation of algorithm (CR-SQN) is the computational effort in calculation of B_k^{-1}. An efficient technique to calculate the inverse is the Cholesky factorization where the matrix B_k is stored in the form of $L_kD_kL_k^T$ and only the matrices L_k and D_k are updated at each iteration. This calculation can be done in $\mathcal{O}(n^2)$ operations (see [13]). In large scale settings, the limited memory variant of the proposed algorithm can be considered which is a subject of our future work.

IV. Numerical experiments

We consider a binary classification problem studied in [23] where the goal is to classify the credit card clients into credible and non-credible based on their payment records and other information. The data set is from the UCI Machine Learning repository. There are 23 features including education, marital status, history of past payment and the mount of bill statement in the past six months. We employ the logistic regression loss function given by (2) where

$$\ell(u_i^Tx, v_i) := -v_i\ln(c(x,u_i)) - (1-v_i)\ln(1-c(x,u_i)) \tag{16}$$

where $c(x,u_i) := (1+\exp(-u_i^Tx))^{-1}$, $v_i \in \{0,1\}$ characterizes the class' type and $u_i \in \mathbb{R}^{23}$ represents the vector of features. We use 1000 data points to run the simulations. We compare the performance of the proposed algorithm (CR-SQN) with that of the regularized stochastic BFGS (RES) algorithm in [17] and also the SA algorithm (SA). To employ RES, since the objective function (2) is non-strongly convex, we assume the function is regularized as in (3) for some constant μ. Fig. 1 and 2 compare the performance the three algorithms. Here we assumed that for CR-SQN, $\rho = 0.9$, $\mu_0 = \delta_k = 1$ for any k, and that γ_k and μ_k are given by (12) with $a = 0.8$, and $c = 0.2$. Also, for RES, we set $\mu = 1$, $\delta = 1$. In both RES and SA schemes, we use $\gamma_k = \gamma_0/(k+1)$. It is observed that in both cases, CR-SQN outperforms RES. Comparing Fig. 1 with Fig. 2, we also observe that the SA scheme seems very sensitive to the choice of the initial stepsize γ_0 which is known as a main drawback of this scheme.

To perform a sensitivity analysis, we compare CR-SQN with RES and SA in Table I and II. In Table I, we report the averaged loss function of CR-SQN and RES for different settings of regularization. We maintain the initial regularization parameter of CR-SQN, μ_0 and the regularization parameter of RES, μ to be equal. We observe that in all settings, CR-SQN attains a lower averaged loss value. In Table II, we observe that by changing the initial stepsize γ_0, except for the case $\gamma_0 = 0.1$, CR-SQN outperforms the SA scheme.

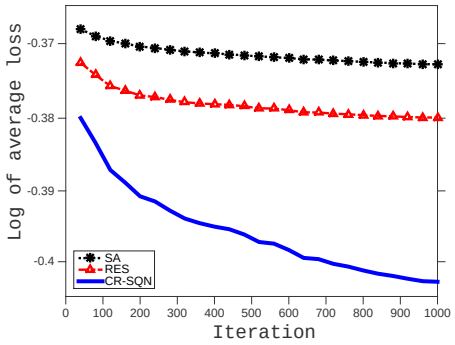

Fig. 1: CR-SQN vs. RES vs. SA - $\gamma_0 = 0.01$

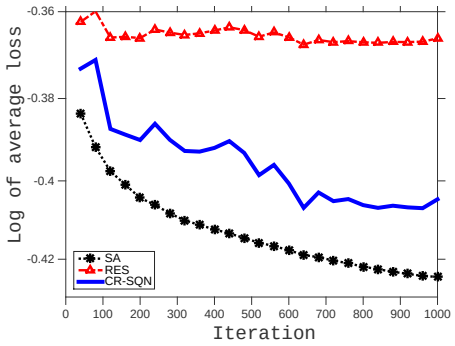

Fig. 2: CR-SQN vs. RES vs. SA - $\gamma_0 = 0.1$

V. Concluding remarks

To address stochastic optimization problems in the absence of strong convexity, we developed a cyclic regularized stochastic SQN method where at each iteration, the gradient mapping and the Hessian approximate matrix are regularized. To maintain the secant condition and carry out the convergence analysis, we do the regularization in a cyclic manner. Under specific update rules for stepsize and regularization parameters, our algorithm generates a sequence that converges to an optimal solution of the original problem in both almost sure and expected senses. Importantly, our scheme is characterized by a derived convergence rate in terms of the objective function values. Our preliminary empirical analysis on a binary classification problem is promising.

CR-SQN		RES	
μ_0	ave. loss	μ	ave. loss
1	0.6684	1	0.6839
0.1	0.6683	0.1	0.6949
0.01	0.8205	0.01	0.9017
0.001	4.3090	0.001	5.0424

TABLE I: CR-SQN vs. RES: varying regularization parameter

CR-SQN		SA	
γ_0	ave. loss	γ_0	ave. loss
0.1	0.6673	0.1	0.6540
0.01	0.6683	0.01	0.6888
0.001	0.6901	0.001	0.6927
0.0001	0.6928	0.0001	0.6931

TABLE II: CR-SQN vs. SA: varying initial stepsize

References

[1] J. C. Spall, *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley, Hoboken, NJ, 2003.

[2] Y. M. Ermoliev, "Stochastic quasigradient methods," in *Numerical Techniques for Stochastic Optimization*. Sringer-Verlag, 1983, pp. 141–185.

[3] V. S. Borkar and S. P. Meyn, "The O.D.E. method for convergence of stochastic approximation and reinforcement learning," *SIAM J. Control Optim.*, vol. 38, no. 2, pp. 447–469 (electronic), 2000.

[4] F. Facchinei and J.-S. Pang, *Finite-dimensional variational inequalities and complementarity problems. Vols. I,II*, ser. Springer Series in Operations Research. New York: Springer-Verlag, 2003.

[5] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of the 19th International Conference on Computational Statistics*, Y. Lechevallier and G. Saporta, Eds. Paris, France: Springer, 2010, pp. 177–187.

[6] R. H. Byrd, S. L. Hansen, J. Nocedal, and Y. Singer, "A stochastic Quasi-Newton method for large-scale optimization," 2015, arXiv:1401.7020v2 [math.OC].

[7] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of Royal Statistical Society*, vol. 58, no. 1, pp. 267–288, 1996.

[8] A. Shapiro, "Monte Carlo sampling methods," in *Handbook in Operations Research and Management Science*. Amsterdam: Elsevier Science, 2003, vol. 10, pp. 353–426.

[9] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statistics*, vol. 22, pp. 400–407, 1951.

[10] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1574–1609, 2009.

[11] F. Yousefian, A. Nedić, and U. V. Shanbhag, "Self-tuned stochastic approximation schemes for non-Lipschitzian stochastic multi-user optimization and nash games," *IEEE Transactions on Automatic Control*, 2015, to appear.

[12] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Math. Program.*, vol. 45, no. 3, pp. 503–528, Dec. 1989. [Online]. Available: http://dx.doi.org/10.1007/BF01589116

[13] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York: Springer, 2006.

[14] D.-H. Li and M. Fukushima, "A modified BFGS method and its global convergence in nonconvex minimization," *Journal of Computational and Applied Mathematics*, vol. 129, pp. 15–35, 2001.

[15] A. Bordes and L. B. N. P. Gallinari, "SGD-QN: Careful quasi-newton stochastic gradient descent," *Journal of Machine Learning Research*, vol. 10, pp. 1737–1754, 2009.

[16] N. N. Schraudolph, J. Yu, and S. Gunter, "A stochastic quasi-newton method for online convex optimization," *In Proc. 11th Intl. Conf. on Artificial Intelligence and Statistics (AIstats)*, pp. 433–440, 2007.

[17] A. Mokhtari and A. Ribeiro, "RES: regularized stochastic BFGS algorithm," *IEEE Transactions on Signal Processing*, vol. 62, no. 23, pp. 6089–6104, 2014.

[18] ——, "Global convergence of online limited memory BFGS," *Journal of Machine Learning Research*, vol. 16, pp. 3151–3181, 2015.

[19] X. Wang, S. Ma, and W. Liu, "Stochastic quasi-Newton methods for nonconvex stochastic optimization," 2014, arXiv:1412.1196 [math.OC].

[20] A. Lucchi, B. McWilliams, and T. Hofmann, "A variance reduced stochastic newton method," *arXiv preprint arXiv:1503.08316 (2015)*.

[21] A. Nedić and S. Lee, "On stochastic subgradient mirror-descent algorithm with weighted averaging," *SIAM Journal on Optimization*, vol. 24, no. 1, pp. 84–107, 2014.

[22] B. Polyak, *Introduction to optimization*. New York: Optimization Software, Inc., 1987.

[23] I. C. Yeh and C. H. Lien, "The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients," *Expert Systems with Applications*, vol. 36, no. 2, pp. 2473–2480, 2007.